Applications of Machine Learning Techniques in Digital Processing of Images of the Martian Surface

Catherine S. Plesko [a*], Steven P. Brumby [a], John C. Armstrong [b], Elliot Ginder [c], and Conway Leovy [d]

[a] Space and Remote Sensing Sciences, Los Alamos National Laboratory,
Mail Stop D436, Los Alamos, New Mexico 87545
[b] Department of Astronomy, University of Washington,
Box 351580, Seattle, Washington 98195
[c] Departments of Mathematics, and Applied and Computational Mathematical Sciences,
University of Washington.
[d] Department of Atmospheric Sciences, University of Washington,
Box 351640, Seattle, Washington 98195

NASA spacecraft have now returned many thousands of images of the surface of Mars. It is no longer practical to analyze such a large dataset by hand, while the development of handwritten feature extraction tools is expensive and laborious. This project investigates the application of machine learning techniques to problems of feature extraction and digital image processing within the Mars dataset.

The Los Alamos GENIE machine learning software system uses a genetic algorithm to assemble feature extraction tools from low-level image operators. Each generated tool is evaluated against training data provided by the user. The best tools in each generation are allowed to "reproduce" to produce the next generation, and the population of tools evolves until it converges to a solution or reaches a level of performance specified by the user. Craters are one of the most scientifically interesting and most numerous features on Mars, and present a wide range of shapes at many spatial scales. We now describe results on development of crater finder algorithms using voting sets of simple classifiers evolved by a machine learning/genetic programming system (the Los Alamos GENIE software).

Keywords: Mars, Meteorite-Craters, Machine Learning, Feature Extraction, Digital Image Processing.

1. INTRODUCTION: CRATER DETECTION WITH THE MARS ORBITAL CAMERA

The Mars Global Surveyor[1] (MGS) satellite has been on orbit around Mars since September 1997. It carries a suite of scientific instruments including the Mars Orbiter Camera[2] (MOC). The MOC is made up of two wide-angle cameras and one narrow-angle, high-resolution camera. The narrow angle dataset, which is the focus of this study, provides imagery with a spatial resolution of the order of 3 meters/pixel and is sensitive to light in a broad visible/near-infrared spectral range (0.50μm – 0.90μm). Since arriving at the planet, the MOC has taken over 93,000 images, which have been used to study a variety of geologic and atmospheric processes including sediments, dust streaks, and volcanism (e.g., see Refs. 1, 3-9). MGS completed its primary mission on January 31, 2001, and continues to take data on areas of interest.

Craters are a rich source of information about planetary surfaces. The number density of craters per unit area as a function of crater diameter can be used to estimate the age of the surface if the impactor flux and erosion rates are known. Alternately, if the surface age is known from other methods, the crater record can be used to examine local erosion processes. Crater surveys are potentially extremely informative, but large amounts of data are difficult to analyze using traditional techniques. The actual counting is currently done by hand, by research students or trained volunteers This technique is limited by the sheer volume of data now being returned by the satellites on orbit around Mars, as well as the inconsistencies in the numbers returned by the counters. For example, seven students counted the craters in MOC image[10] M1500956 for this study. The students returned a different total for each survey, even if the same student was recounting the image, as shown in Table 1. The average number of craters found was 222. We would describe the magnitude of variability in the returned results as disappointing.

* Further author information: (Send correspondence to C.S.P.) Email: plesko@lanl.gov

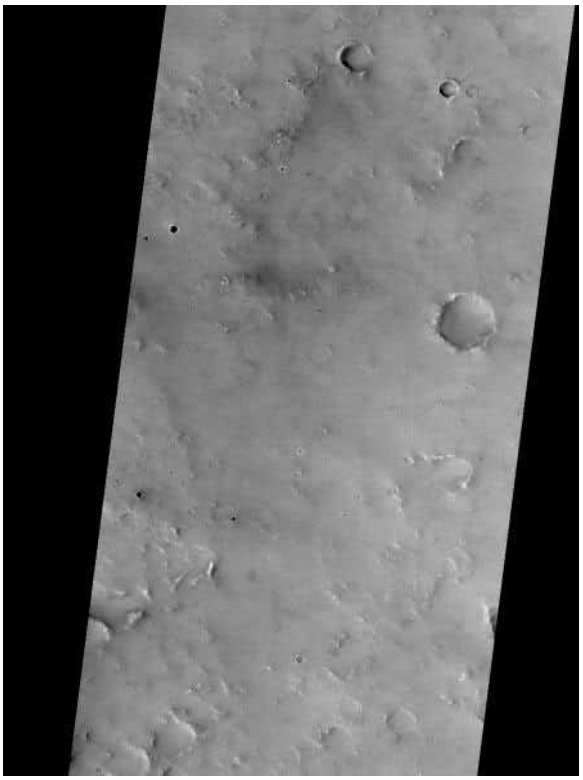

Figure 1. The first thousand pixel rows of one of the 93,000 images returned by the Mars Global Surveyor Mars Orbital Camera (MGS/MOC). Image M1500956 (see Ref. 10) is of an inter-crater plane in the Deucalionis Regio area of the Martian southern hemisphere. Detection and counting of small young craters usual traditional manual techniques by trained graduate students and volunteers generally produces a highly uncertain estimate for any given scene.

Hand count accuracy can be improved by averaging the results of many individual counts, but doing so multiplies an already ponderous dataset by factors of ten or even hundreds. A preliminary study of this sort was conducted by Kanefsky, et. al.[11], at NASA Ames Research Center, using untrained volunteers to hand mark craters through a web-page interface. They found that averaging the results of two hundred volunteers closely matched the results of an experienced image analyst on a given image. A similar technique can be used with automated classification algorithms. One of the simplest commonly used approaches to combining classifiers is through a process of voting, in which a group of binary classifiers (i.e., the answer produced by the classifier is "yes" for feature or "no" for non-feature) vote on each data point, and the majority decision holds. To be successful, this approach assumes that the mistakes made by different classifiers tend to be uncorrelated across classifiers, while detections of the feature of interest are correlated. For example, separate classifiers that trigger on "smoothness" and "brightness" might be successfully combined to detect young, uncratered surfaces. Voting can be useful

student	first count	second count
1	87	
2	104	
3	273	
4	236	304
5	436	
6	222	
7	113	

Table 1. Comparison of hand-counts of craters in Figure 1.

when it is "cheap" to learn reasonably accurate classifiers (in terms of training time or computational effort) but expensive to obtain a single, highly accurate classifier.

For the present study we describe the development of voting sets of crater finding algorithms, using a genetic programming approach originally developed for Earth remote sensing applications. We are initially interested in developing pixel-level algorithms that segment out spatial regions likely to contain craters, and in future work will explore higher-level, scale-invariant processing which uses the geometry of the segmented regions (e.g., linear feature, circular feature, annular feature) to complete the classification of craters as geometric objects. We present initial results with sets of classifiers, combined using a voting procedure to improve performance.

2. GENIE: GENETIC PROGRAMMING FOR AUTOMATIC FEATURE EXTRACTION

Los Alamos National Laboratory's GENIE software[12-15] is a machine learning software system using techniques from the fields of genetic algorithms (GA)[16-18] and genetic programming (GP)[19] to construct spatio-spectral feature extraction algorithms for multi-spectral remotely sensed imagery. Both the structure of the feature extraction algorithm, and the parameters of the individual image processing steps, are learned by the system. GENIE has been described at length elsewhere[12-15], so we will only present a brief description of the system here. In particular, the present work explores using GENIE on panchromatic imagery.

GENIE follows the paradigm of genetic programming: a population of candidate image-processing algorithms is randomly generated from a collection of low-level image processing operators, including texture measures, spectral band-math operations (e.g. ratios of bands), and various morphological filters. The fitness of each individual is assessed from its performance on training data provided by the human user via a graphical interface. Our fitness metric in based on measuring the total error rate (false positives and false negatives) on the feature extraction task. After a fitness value has been assigned to each candidate algorithm in the population, the most fit members of the population reproduce with modification via the evolutionary operators of mutation and crossover. This process of fitness evaluation and reproduction with modification is iterated until the population converges, or some desired level of classification performance is attained, or some user-specified limit on computational effort is reached (e.g. number of candidate algorithms evaluated). The final result is a gray-scale enhancement of the feature of interest, which is then converted into a final Boolean classification using a threshold.

The algorithms evolved by GENIE combine spatial and spectral processing, and the system was designed to enable exploration of spatio-spectral image processing. This system has been shown to be effective in detecting complex spatio-spectral terrain features in multispectral imagery, such as golf courses in MODIS Airborne Simulator imagery[20], and in a range of real world problems, including delineating and classifying wildfire burn scars[21] and vegetation land-cover classes[22] using a number of multispectral imagery datasets, and earlier work by the present authors on proposing the detection of craters on Mars[23] using a high-resolution panchromatic dataset (Mars Global Surveyor/Mars Orbital Camera). In that work, individual classifiers were evolved and encouraging qualitative results obtained, but performance outside the training area showed a substantial degradation of performance. This led us to take interest in schemes for suppressing false alarms, including the voting scheme reported in the present work.

We now briefly describe our method of providing training data, our encoding of image-processing algorithms as chromosomes for manipulation by the GA, our method for evaluating the fitness of individuals in the population, and the voting scheme adopted.

2.1 Training Data

The environment for the population consists of one or a number of training scenes. Each training scene contains a raw image, together with a weight plane and a truth plane. The weight plane identifies the pixels to be used in training, and the truth plane locates the features of interest in the training data. Providing sufficient quantities of good training data is crucial to the success of any machine learning technique. In principle, the weight and truth planes may be derived from an actual ground campaign (i.e., collected on the ground at the time the image was taken), may be the result of applying some existing algorithm, and/or may be marked-up by hand using the best judgement of an analyst looking at the data. We have developed a graphical user interface (GUI), called Aladdin, for manual marking up of raw imagery. Using Aladdin, the analyst can view a raw image in a variety of ways, and can mark up training data by painting directly on the image using the mouse. Training data is ternary-valued, with the possible values being "true", "false", and "unknown". True defines areas where the analyst is

confident that the feature of interest does exist. False defines areas where the analyst is confident that the feature of interest does not exist. Unknown pixels do not influence the fitness of a candidate algorithm.

2.2 Representation of Image-Processing Algorithms

Traditional genetic programming[19] (GP) uses a variable sized (within limits) tree representation for algorithms. Our representation differs in that it allows for reuse of values computed by sub-trees, i.e. the resulting algorithm is a graph rather than a tree. The image processing algorithm that a given chromosome represents can be thought of as a directed acyclic graph where the non-terminal nodes are primitive image processing operations, and the terminal nodes are individual image planes extracted from the multi-spectral image used as input. Our representation also differs in that the total number of nodes is fixed (although not all of these may actually be used in the final graph), and crossover is carried out directly on the linear representation.

We have restricted our "gene pool" to a set of useful primitive image processing operators ("genes"). These include spectral, spatial, logical and thresholding operators. The set of morphological operators is restricted to function-set processing morphological operators, i.e., gray-scale morphological operators having a flat structuring element. The sizes and shapes of the structuring elements used by these operators is also restricted to a pre-defined set of primitive shapes, which includes the square, circle, diamond, horizontal cross and diagonal cross, and horizontal, diagonal, and vertical lines. The shape and size of the structuring element are defined by operator parameters. Other local neighborhood/windowing operators such as mean, median, etc., specify their kernels/windows in a similar way. We define scratch planes as blocks of memory for storing intermediate calculations within a candidate image-processing algorithm. Once "scratch" planes have been generated, GENIE is allowed to explore weighted sums, differences and ratios of data and scratch planes.

A single gene consists of an operator, plus a variable number of input arguments specifying from where input is read, output arguments specifying where output is to be written, and any additional parameters that might be required to specify how the specific operator works (e.g., the diameter and shape of a structuring element used in a morphological filter). The operators used in Genie take one or more distinct image planes as input, and generally produce a single image plane as output. Input can be taken from any data plane in the training data image cube. Output is written to one of a number of scratch planes, temporary workspaces where an image plane can be stored. Genes can also take input from scratch planes, but only if that scratch plane has been written to by another gene positioned earlier in the chromosome sequence. We use a notation for genes[12,14] that is most easily illustrated by an example: the gene `[ADDP rD0 rS1 wS2]` applies pixel-by-pixel addition to two input planes, read from data plane 0 and from scratch plane 1, and writes its output to scratch plane 2. Any additional required operator parameters are listed after the output arguments.

Note that although all chromosomes have the same fixed number of genes, the effective length of the resulting algorithm may be smaller than this. For instance, an operator may write to a scratch plane that is then overwritten by another gene before anything reads from it. GENIE performs an analysis of chromosome graphs when they are created and only carries out those processing steps that actually affect the final result. Therefore, the fixed length of the chromosome acts as a maximum effective length.

2.3 Supervised Classification and Fitness Evaluation

Each candidate image-processing algorithm generates a number of intermediate feature planes (or "signature" planes), which are then combined to generate a Boolean-valued mask for the feature of interest. This combination is achieved using a standard supervised classifier (we use the Fisher linear discriminant[24]), and an optimal threshold function.

Complete classification requires that the image-processing algorithm produce a binary-valued output plane for any given scene. It is possible to treat, e.g., the contents of the first scratch plane as the final output for that candidate image- processing algorithm (thresholding would generally be required to obtain a binary result, though Genie can choose to apply its own Boolean thresholding functions). However, we have found it to be useful to perform the combination of the data and scratch planes using a non-evolutionary method, and have implemented a supervised classifier backend. To do this, we first select a subset of the scratch planes and data planes to be "signature" planes. For the present experiments, this subset consists of just the scratch planes. We then use the provided training data and the contents of the signature planes to derive the Fisher Discriminant, which is the linear combination of the signature planes that maximizes the mean separation in spectral terms between those pixels marked up as "true "and those pixels marked up as "false", normalized by the total variance in the projection defined by the linear combination. The output of the discriminant-finding phase is a real-valued single-plane

Class.	Training Scene			Test Scene 1			Test Scene 2		
	Score	D.R.	F.A.R.	Score	D.R.	F.A.R.	Score	D.R.	F.A.R.
1	915.1	90.65	7.63	869.1	91.04	17.22	774.9	98.82	43.84
2	969.3	97.56	3.7	959.1	94.78	2.96	950.9	98.40	8.22
3	950.5	96.29	6.19	897.1	95.34	15.93	761.1	99.39	47.17
4	964.9	96.44	3.46	954.2	93.48	2.64	954.1	98.38	7.56
5	979.7	98.94	3.00	960.1	94.76	2.74	951.3	98.03	7.77
6	949.6	94.26	4.34	908.4	92.09	10.41	796.9	99.47	40.09
Vote 1	980.3	97.31	1.24	960.0	94.33	2.34	950.2	98.50	8.46
Vote 2	972.6	97.48	2.97	935.9	94.77	7.59	818.9	99.51	35.73
Vote 3	923.4	84.83	0.15	919.3	84.38	0.52	966.1	96.77	3.54

Table 2. Results of individual classifiers and of voting. Six individual classifiers ("Class.") were evolved using GENIE, and then combined using voting with a majority rule. Classifier "Vote 1" gives the results achieved by a majority voting scheme with contributions from all six classifiers. Classifier "Vote 2" gives the results achieved by a majority voting scheme with contributions from all classifiers apart from the strongest individual classifier, Classifier 5. The "Score" is the score reported by GENIE, in the range 0 – 1000 where 1000 is a perfect score. D.R. is the detection rate in percent, and F.A.R. is the false alarm rate in percent. Classifier "Vote 3" shows the result of adopting a "unanimous" voting decision rule for the all-component "Vote 1" classifier.

"answer" image. This is reduced to a binary image by exhaustive search over all the training pixels to find the threshold value that minimizes the total number of misclassifications (false positives plus false negatives) on the training data.

The fitness of a candidate solution is given by the degree of agreement between the final binary output plane and the training data. This degree of agreement is determined by the Hamming distance between the final binary output of the algorithm and the training data, with only pixels marked as true or false (as recorded in the weight plane) contributing towards the metric. The Hamming distance is then normalized so that a perfect score is 1000.

2.4 Combining classifiers using voting

We adopt a simple voting scheme: a set of reasonably fit binary classifiers is evolved in a series of separate GENIE runs. GENIE's genetic algorithm is a stochastic learning process, so individual results are likely to be highly variable in structure. The binary classification outputs of these classifiers is then summed and thresholded at whatever level we choose to define the joint decision. In general, a simple majority rule is adopted, though if the classifiers are noisy, a unanimous rule decision may also be adopted. Voting classifier schemes have been in use for a long time (see, e.g., the general review in Ref. 25), and have led to more elaborate schemes for building a sequence of supervised classifiers trained on specially selected subsets of the available labeled training data. The extension of these more complex schemes to highly correlated datasets, such as the pixels of a high-resolution image, is currently a topic of research.

3. RESULTS

We selected a training image (MOC image M0803054, near Louros Valles (8.5S, 82.0W), Ref. 26), chosen to present a reasonably homogeneous terrain marked by a number of bowl shaped craters obvious to the human eye. GENIE was trained on the first 930 pixel rows of the image (the image is 830 pixels wide) with a truth file based on manual analysis (Fig. 2; result with voting shown in Fig. 3) in which the analyst has marked some of the fresh, bowl shaped craters in the scene as true, and some of the protruding surface features and non-cratered terrain as false. The next 970 pixel rows of this scene where also marked by hand, and kept back to serve as our Test Scene 1 (Fig. 4). The final 530 pixel rows of this image show a noticeable brightening due to angle of Solar illumination (Fig. 5), and hence this part of the scene presents a tougher classification challenge for any classifier trained only on the first 930 pixel rows of the scene. We use this final 530 pixel row image as our Test Scene 2.

GENIE was run 6 times, each time with a new population of 30 chromosomes per generation, each run lasting for 50 generations. Running on standard Intel/Linux workstations, each run required approximately 1 hour of wall-clock time. Each

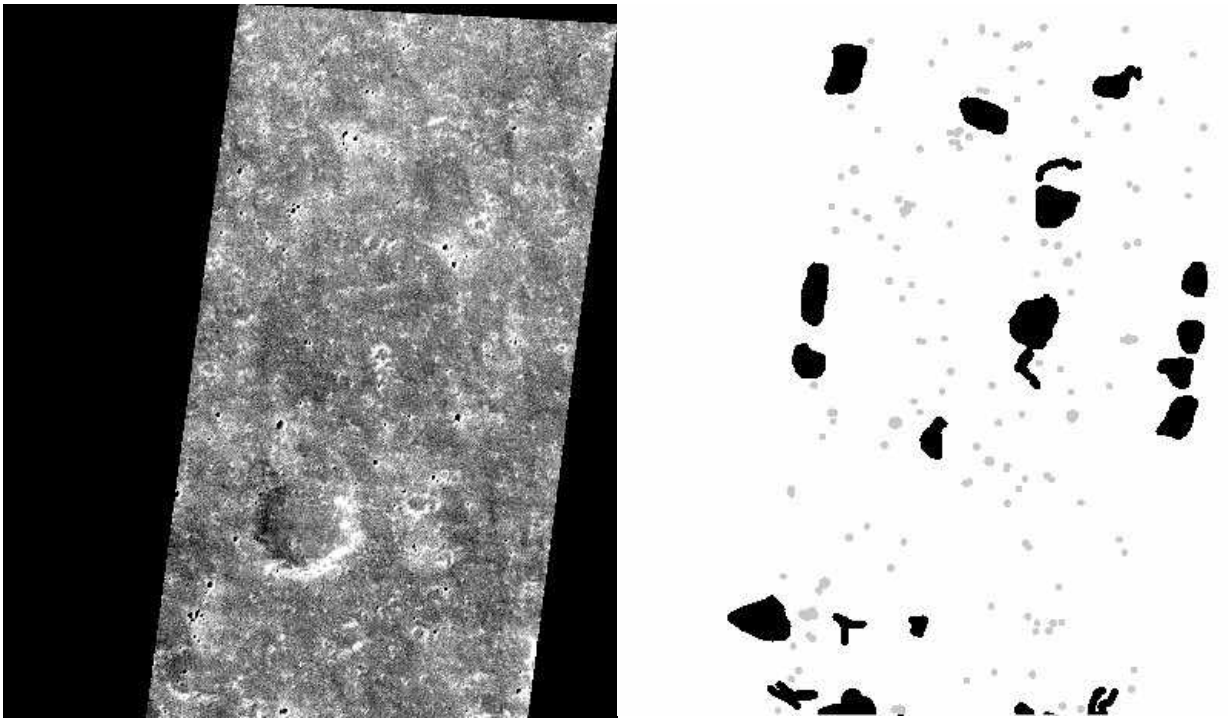

Figure 2. MOC Image M0803054: Training Scene and Training Data. The training scene shown (Left) is an 830x930 pixel region imaged by the MOC high-resolution instrument. The training data (Right), obtained by traditional inspection of the image by a trained image analyst graduate student, shows the locations of regions containing young bowl-shaped craters (marked in light gray) and locations of example non-crater features (marked in black).

candidate algorithm contained a maximum of 18 image processing operations, and used three intermediate scratch planes. Results for individual classifiers and for the combined "voting" classifier, is shown in Table 2.

The best individual crater finding algorithm achieved a training score of 980, with a detection rate of 97% and a false alarm rate of 3%. On the test data, the performance of this classifier dropped as described in Table 2) to a score of 960 on the similar Test Scene 1, and to 951 on the noticeably differently illuminated The algorithm found 9 image processing steps, comprised of a mixture of morphological spatial filters and texture and neighborhood statistics measures. Note that even though we specified a maximum of 18 image processing for the candidate algorithms, only 9 image processing steps appear in the final best algorithm. In the language of genetic algorithms, the remaining 9 potential image processing steps remain "unexpressed genes", in principle providing robustness against mutation and providing room for further adaptation if we later decide to update the training data and continue the evolution of candidate solutions.

Classifier "Vote 1" gives the results achieved by a majority voting scheme with contributions from all six classifiers. This combined classifier slightly out-performs the best individual classifier on the training scene, and performs slightly worse than the single best classifier when applied out of sample. This combined classifier out-performs all the other classifiers on the training scene and on Test Scene 1, and has average performance on Test scene 2. The situation on Test Scene 2 is complex, and the small amount of data is suggestive of a split in the behavior of individual classifiers, between those that continue to work under the changed illumination conditions, and those that fail under previously unexperienced conditions. As expected, the false alarm rate reported by the voting set of classifiers is substantially lower than the false alarm rate reported by any other classifier on the training scene or on Test Scene 1, and achieves results comparable to the best individual classifier on Test Scene 2.

Results with classifier "Vote 1" raise the question, how much of the behavior of this classifier is simply due to the good individual Classifier 5? Classifier "Vote 2" gives the results achieved by a majority voting scheme with contributions from all classifiers apart from the strongest individual classifier, Classifier 5. As can be seen, detection rates are maintained relative to classifier "Vote 1", while false alarm rates rise somewhat.

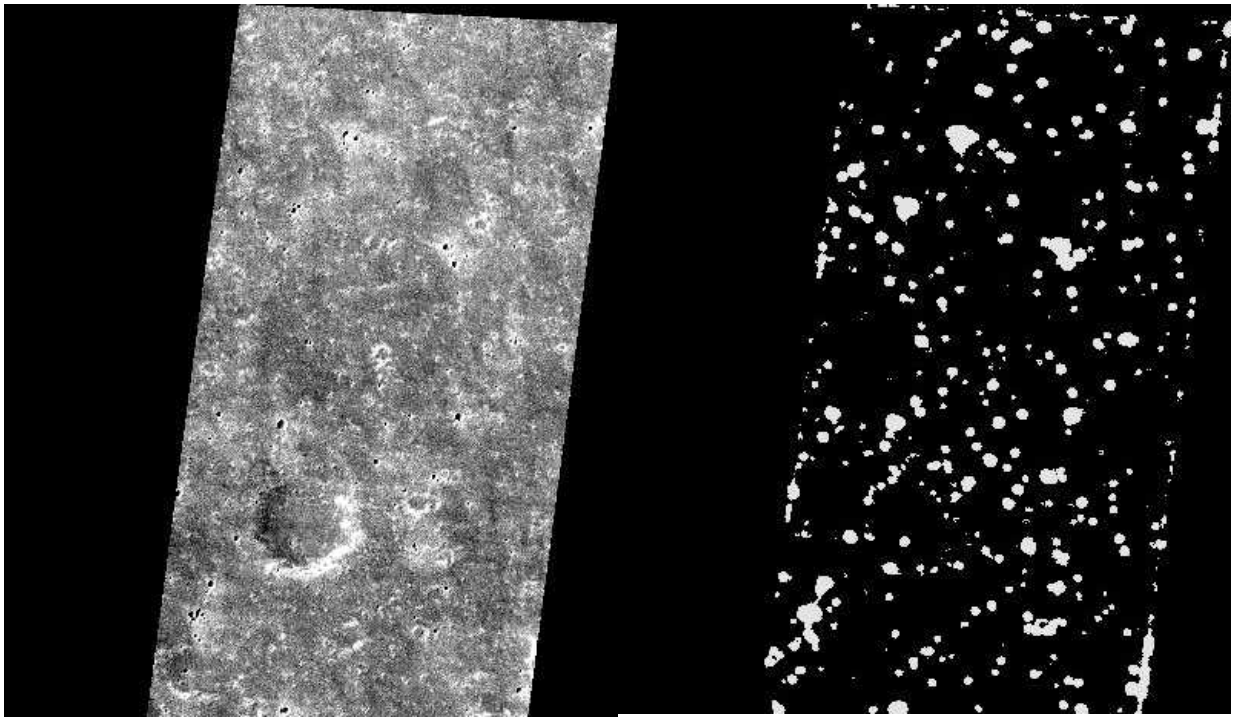

Figure 3. MOC Image M0803054: Training Scene and Result. Result (Right) obtained by majority vote of six independently evolved GENIE image processing classifiers for regions containing young bowl-shaped craters. See Table 2 for quantitative analysis.

Classifier "Vote 3" shows the result of adopting a "unanimous" voting decision rule for the all-component "Vote 1" classifier. For application where achieving low false alarm rate is paramount, this is clearly a superior classifier. The extent to which this low-noise classifier can be used as the initial stage of a geometric-based extraction of craters as objects, will be the subject of future work.

4. CONCLUSIONS

The surface of Mars has been observed from orbit at a number of progressively finer spatial scales, from kilometers (Mariner series), to 100's of meters (Viking), and most recently at the few meter scale (Mars Global Surveyor, and Mars Odyssey). These new data sets hold out the greatest chance of significant progress in understanding the history of the Martian surface, while also presenting the greatest challenge to traditional manual analysis techniques. This study investigated the evolution of a voting set of crater finding algorithms for application to the Mars Orbiter Camera narrow angle dataset. We described results on training and test images. The algorithm is successful at detecting craters within the images, and generalized well to an image that it had not seen before. We find these results to be encouraging for the application of GENIE to the MOC panchromatic dataset.

ACKNOWLEDGEMENTS

C.S.P. wishes to thank Michael Gillis, Marilee Glenn, Michelle Hansard, and Marty Smyth, for their assistance in manual crater counts. C.S.P. also acknowledges the support of a Mary Gates Foundation research grant, and the financial support of the Washington NASA Space Grant Consortium. S.P.B. would like to thank his colleagues on the GENIE software team: J.J. Bloch, M. Galassi, N.R. Harvey, S. Perkins, R.B. Porter, J.J. Szymanski, J. Theiler, and A.C. Young. GENIE research and development is funded by the U.S. Departments of Defense and Energy.

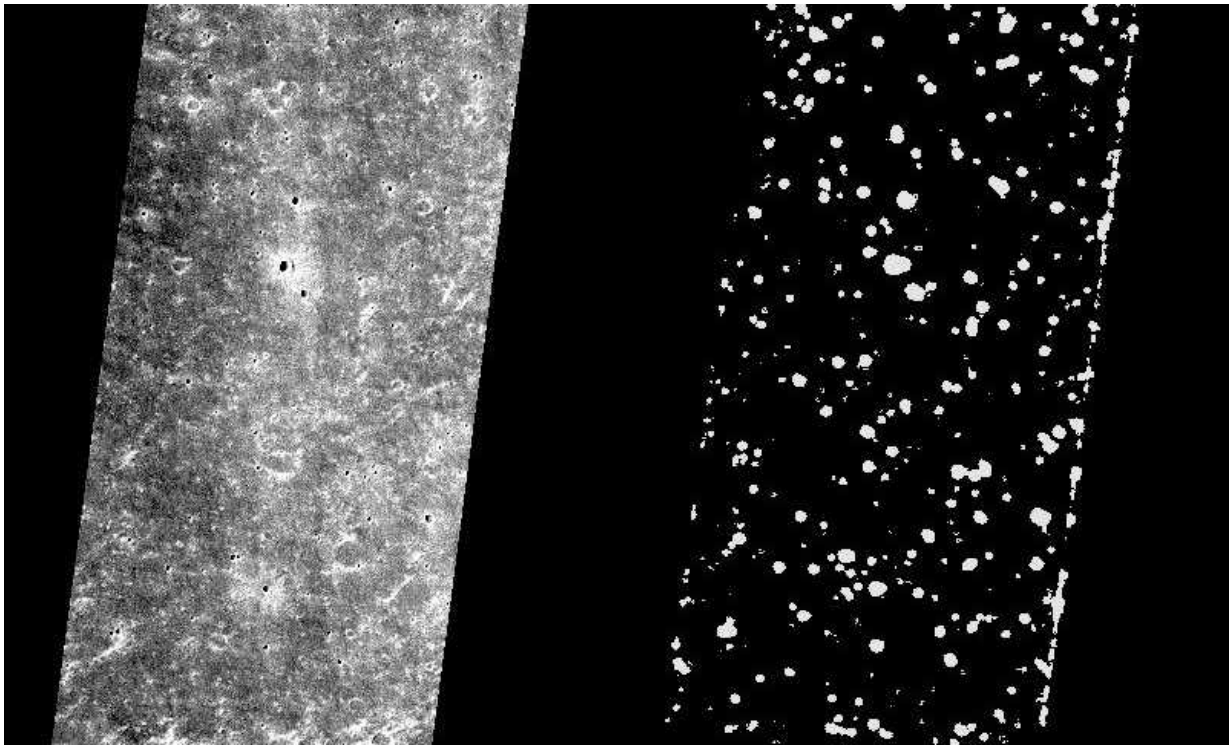

Figure 4. MOC Image M0803054: Test Scene 1 and Result. Result (Right) obtained by majority vote of six independently evolved GENIE image processing classifiers for regions containing young bowl-shaped craters. See Table 2 for quantitative analysis. Test Scene 1 is similar in character to the training scene.

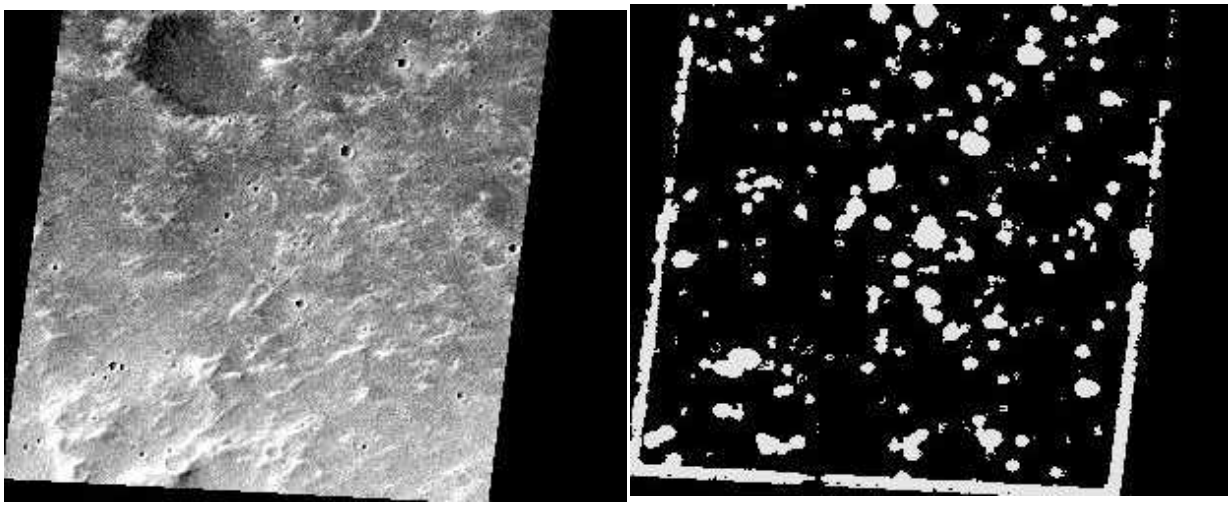

Figure 5. MOC Image M0803054: Test Scene 1 and Result. Result (Right) obtained by majority vote of six independently evolved GENIE image processing classifiers for regions containing young bowl-shaped craters. See Table 2 for quantitative analysis. Test Scene 2 has a noticeably different general illumination level at the southern edge of the image, which presents a stronger test of generalization of the classifiers evolved by GENIE.

REFERENCES

1. A. L. Albee, F. D. Palluconi, R. E. Arvidson, “Mars Global Surveyor Mission: Overview and Status”, Science, Vol. 279, p. 13, 1998.
2. M. C. Malin, G. E. Danielson, A. P. Ingersoll, H. Masursky, J. Veverka, M. A. Ravine, T. A. Soulanille, "Mars Observer Camera", Journal of Geophysical Research, Vol. 97(E5), pp. 7699-7718, 1992.
3. K. S. Edgett, and M. C. Malin, “New views of Mars eolian activity, materials, and surface properties: Three vignettes from the Mars Global Surveyor Mars Orbiter Camera,” Journal of Geophysical Research, Vol. 105(E1), pp.1623-1650, 2000.
4. W. K. Hartmann, M. Malin, A. McEwen, M. Carr, L. Soderblom, P. Thomas, E. Danielson, P. James, and J. Veverka, “Evidence for recent volcanism on Mars from crater counts,” Nature, Vol. 397, pp. 586-589, 1999.
5. P. C. Thomas, M. C. Malin, M. H. Carr, G. E. Danielson, M. E. Davies, W. K. Hartmann, A. P. Ingersoll, P. B. James, A. S. McEwen, L. A. Soderblom, and J. Veverka, “Bright dunes on Mars,” Nature, Vol. 397, pp. 592-594, 1999
6. J.R. Zimbelman, "Non-active dunes in the Acheron Fossae region of Mars between the Viking and Mars Global Surveyor eras", Geophysical Research Letters. Vol.27, no.7, pp.1069-1072, 2000.
7. M. H. Carr, M. C. Malin, "Meter-scale characteristics of Martian channels and valleys", Icarus., Vol.146, no.2, p.366-386, 2000.
8. Hartmann, W.K.,"Martian cratering. VI. Crater count isochrons and evidence for recent volcanism from Mars Global Surveyor", Meteoritics and Planetary Science. vol.34, no.2, p.167-77, 1999.
9. Hartmann, W.K., Berman, D.C., "Elysium Planitia lava flows: crater count chronology and geological implications", Journal of Geophysical Research. vol.105, no.E6, p.15011-25, 2000.
10. M.C. Malin, K. S. Edgett, M. H. Carr, G. E. Danielson, M. E. Davies, W. K. Hartmann, A. P. Ingersoll, P. B. James, H. Masursky, A. S. McEwen, L. A. Soderblom, P. Thomas, J. Veverka, M. A. Caplinger, M. A. Ravine, T. A. Soulanille, and J. L. Warren, “Intercrater plain in Deucalionis Regio”, NASA's Planetary Photojournal (http://photojournal.jpl.nasa.gov/), Catalog number M1500956, Created 5-14-2000.
11. Kanefsky, B., N. Barlow, and V. Gulick “The Clickworkers Study: NASA Ames’s Experiment in Volunteer Science.” San Jose: NASA Ames Research Center. http://clickworkers.arc.nasa.gov/top accessed June 5, 2002.
12. S.P. Brumby, et al., “Investigation of feature extraction by a genetic algorithm ”, Proc. SPIE, Vol. 3812, pp. 24-31,1999; see also the GENIE web site http://www.daps.lanl.gov/genie .
13. J. Theiler, et al., “Evolving retrieval algorithms with a genetic programming scheme ”, Proc. SPIE, Vol. 3753, pp.416-425,1999.
14. N.R. Harvey, et al., “Parallel evolution of image processing tools for multispectral imagery”, Proc. SPIE, Vol. 4132, pp. 72-82, 2000.
15. S. Perkins, et al., “GENIE: A hybrid genetic algorithm for feature classification in multi-spectral images”, Proc. SPIE, Vol. 4120, pp 52-62, 2000.
16. J. H. Holland, Adaptation in Natural and Artificial Systems, University of Michigan, Ann Arbor (1975).
17. I. Rechenberg, Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution, Fromman-Holzboog, Stuttgart (1973).
18. L. Fogel, A. Owens and M. Walsh, Artificial Intelligence through Simulated Evolution, Wiley, New York (1966).
19. J. R. Koza, Genetic Programming: On the Programming of Computers by Natural Selection, MIT, Cambridge (1992).
20. N. R. Harvey, S. Perkins, S. P. Brumby, J. Theiler, R. B. Porter, A. C. Young, A. K. Varghese, J .J. Szymanski, and J. J. Bloch, *Finding golf courses: The ultra high tech approach*, Proc. Second European Workshop on Evolutionary Computation in Image Analysis and Signal Processing (EvoIASP2000), Edinburgh, UK, pp. 54-64, 2000.
21. S. P. Brumby, N. R. Harvey, J. J. Bloch, J. Theiler, S. Perkins, A. C. Young, and J. J. Szymanski, *Evolving forest fire burn severity classification algorithms for multi-spectral imagery*, Proc. SPIE, Vol. 4381, pp. 236-245, 2001.
22. S. P. Brumby, J. Theiler, J. J. Bloch, N. R. Harvey, S. Perkins, J. J. Szymanski, and A. C. Young, Evolving land cover classification algorithms for multi-spectral and multi-temporal imagery, Proc. SPIE, Vol. 4480, pp. 120-129, 2002.
23. Catherine S. Plesko, Steven P. Brumby and Conway Leovy, Automatic Feature Extraction for Panchromatic Mars Global Surveyor Mars Orbiter Camera Imagery, Proc. SPIE, Vol. 4480, pp. 139-146, 2002.
24. For example, see C.M. Bishop, Neural Networks for Pattern Recognition, pp.105 –112, Oxford University (1995).
25. R. O. Duda, P. E. Hart, and D. G. Stork, Pattern Classification, 2nd ed., Wiley-Interscience, New York (2001).
26. M.C. Malin, K. S. Edgett, M. H. Carr, G. E. Danielson, M. E. Davies, W. K. Hartmann, A. P. Ingersoll, P. B. James, H. Masursky, A. S. McEwen, L. A. Soderblom, P. Thomas, J. Veverka, M. A. Caplinger, M. A. Ravine, T. A. Soulanille, and J. L. Warren, "MOC narrow-angle image M08-03054: Sample", NASA's Planetary Photojournal (http://photojournal.jpl.nasa.gov/), Catalog number M0803054, Created 10-14-1999.